DOCUMENT RESUME

ED 463 307 TM 033 739

AUTHOR Cromwell, Susan
TITLE A Primer on Ways To Explore Item Bias.
PUB DATE 2002-02-00
NOTE 23p.; Paper presented at the Annual Meeting of the Southwest Educational Research Association (25th, Austin, TX, February 14-16, 2002).
PUB TYPE Information Analyses (070) -- Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Chi Square; *Item Bias; Item Response Theory; Scores
IDENTIFIERS Delta Plot Method; *Item Bias Detection

ABSTRACT

This paper provides an introductory primer on methods for exploring item bias. The purposes of item bias analysis are to investigate whether test scores are affected by different sources of variance in the various subpopulations, and if different sources of variance are found, to determine if an unfair advantage exists. The review discusses three primary methods for detecting item bias: (1) methods based on item response theory (IRT); (2) chi-square methods; and (3) methods based on delta plots. When choosing the best item bias detection methodology, the researcher must consider the application of the results, the availability of software needed, and the practicality of implementing the methodology chosen. The IRT model leads the other item bias methods by being the most theoretically sound, although statistically complex, procedure. Chi-square and delta plot methods are not as theoretically sophisticated yet are more practical and easier to implement. (Contains 3 tables, 3 figures, and 15 references.) (SLD)

Reproductions supplied by EDRS are the best that can be made
from the original document.

Running head: Primer on Item Bias Detection

ED 463 307

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

S. Cromwell

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

1

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

✓ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

A Primer on Ways to explore Item Bias

Susan Cromwell

Texas A&M University

TM033739

Paper presented at the annual meeting of the Southwestern Educational Research Association, Austin, TX, February 14-15, 2002.

BEST COPY AVAILABLE

Abstract

As Crocker and Algina (1986) noted, "Because irrelevant sources of [score] variation are unavoidable, it is important that they should not give an unfair advantage to one subpopulation of examinees… over another subpopulation" (p. 376). The effort to identify such influences is the object of item bias analyses. The purpose of the present paper is to provide an introductory primer on methods for exploring item bias. The review will discuss three primary methods for detecting item bias: methods based on Item Response Theory, chi-square methods, and methods based on delta plots.

A Primer on Ways to Explore Item Bias

Test scores are inevitably affected by sources of variation other than the construct allegedly measured by the test. If tests invariably measured only what the researcher wanted the tests to measure all scores would be perfectly reliable and valid. However, irrelevant sources of variation cannot be completely controlled; therefore, steps should be taken to avoid giving any unfair advantages to the subpopulations taking the test. Unfair advantages could arise in subpopulations (e.g., females versus males) when compared to another subpopulation. This unfair advantage will exist if within the subpopulations both have equal standing on the construct of interest, yet the irrelevant sources of variation are differentially distributed for the two subpopulations. For example, males and females were asked to read several articles on learning how to sail, and were then given a test on their expected sailing ability. Presume the females grew up near the ocean and sailed extensively, while the males had never been on a sailboat. The knowledge the females gained through the actual experiences of sailing may cause an unfair advantage when taking the test (Crocker & Algina, 1986).

The identification of item bias is a daunting challenge. Systematic differences in item performance for certain subpopulations may be due to item content, such as invoking vocabulary typically known only by a given subpopulation. Jenson (1980) gives an excellent example of item content such as vocabulary with Robert L. William's Black Intelligence Test of Cultural Homogeneity (BITCH-100), which was composed entirely of words, terms, and expressions peculiar to the black culture. On the other hand, differential item performance may instead be due not to the item content itself, but rather to differential access to educational opportunity. For example, the experiences the

females had in sailing that the males did not have in the above example became an unfair advantage. Item differences in this case reflect real subpopulation differences not due to the item content, and the differences would dissipate if the differences in educational opportunity were remediated. This, of course, is easier said than done. Once the differences in sources of variation are detected, determining why subpopulations perform differently on selected items is not an easy task!

Crocker and Algina (1986) stated that the two purposes of item bias analyses are (1) to investigate whether test scores are affected by different sources of variance in the various subpopulations, and if different sources of variance are found, then (2) to determine if an unfair advantage exists. Based on the two purposes of item bias, Crocker and Algina defined a set of items as unbiased if:

> (1) the items are affected by the same sources of variance in both subpopulations; and
>
> (2) among examinees who are at the same level on the construct purportedly measured by the test, the distributions of irrelevant sources of variation are the same for both subpopulations.

Purpose of the Paper

The purpose of the present paper is to provide an introductory primer on methods for exploring item bias. Jenson (1980) provided a comprehensive summary of these techniques, while Henson (1999) and Fisk (1991) provided summaries of methods of item bias as well. This review will discuss three primary methods for detecting item bias: methods based on Item Response Theory, chi-square methods, and methods based on

delta plots. Because delta plots are graphical and the least technical, this procedure will be explained in some detail.

Methods of Item Bias

Item Response Theory Overview

The most theoretically sound method of evaluating item bias invokes the Item Response Theory (IRT), also known as the latent trait method (Fan, 1998; McKinley & Mills, 1989). IRT is based on two postulates that state (1) performance of an examinee on a test item can be explained by a set of factors called latent traits or abilities, and (2) the relationship between examinees' item performance and the traits underlying item performance can be described by an item characteristic curve (ICC) (Henard, 2000). Estimates of item parameters, such as ability level (latent trait) and the probability of responding correctly are analyzed on ICCs to detect item bias. Basically, subpopulation ICCs are compared to one another to determine whether there is item bias for each item on the test.

Before the investigation for item bias begins, the researcher must choose the IRT model to be used. There are several models to choose from in IRT, yet the most commonly used models for item bias are the one-parameter, two-parameter, and three-parameter models. The one-parameter model, also known as the Rasch model, estimates item difficulty, or the b parameter. This is the most parsimonious of the three models. The two-parameter model obtains a b value and an item discrimination statistic, the a parameter. The three-parameter model allows for an additional parameter known as the c parameter, or guessing parameter, to be estimated (Henson, 1999). For a comprehensive review of the three different parameter models see Henson (1999). The most

theoretically sound and statistically complex procedure for measuring item bias is the three-parameter model, which will be illustrated below. The three-parameter model contains all three parameters estimated by IRT, which leads to this model requiring the highest degree of statistical sophistication, a costly program (LOGIST or a similar program), and of course an infamously large sample size (Fisk, 1991; Ironson, 1982). All three parameter models will be illustrated (Fisk, 1991).

IRT with Item Bias

When using IRT to investigate item bias a set of items is deemed unbiased if the ICCs for every item are the same for both subpopulations. All figures and tables for IRT examples are found in Fisk (1991). Before comparing the ICCs, two steps must be taken. First, estimates of the item parameters are calculated and then expressed on the same scale. Scaling parameters is a complex process, and a more in depth explanation can be found in Crocker and Algina (1986) and Henard (2000). Table 1 illustrates scaled item parameters of data for males vs. females that detect item bias from Crocker and Algina (1986). Notably, the items to worry about are item 1 and 5, and possibly 3. These three items may need to be rewritten to remove the source of bias. Item 8 is your lowest, and therefore the least biased item.

Insert Table 1 about here

Second, an index of item bias, also known as the estimate of bias, is calculated for each item by calculating the area between ICCs. The estimate of bias may be calculated in several different ways depending on the parameter model chosen (Fisk, 1991). The three-parameter model and the Rasch model (one-parameter) will be illustrated.

Three-Parameter Model

The Reimann sum approximation and proposed null hypothesis by Lord (1980) is used for the three-parameter model. Formulas for both methods are listed below.

Reimann sum approximation:

$$A_g = S(b-a)/n \, l \, P_{1g}[a+k(b-a)/n] - P_{2g}[a+k(b-a)/n]$$

Lord's (1980) alternative hypothesis for Null:

$$\text{Ho: } a_{1g} = a_{2g}, b_{1g} = b_{2g}$$

For the Reimann sum approximation, the estimate of bias (A_g) is yielded by the following formula. This formula yields the approximate area of A_g if the probabilities P_{ig} (q) are given on some interval $a \leq q \leq b$ in increments of (b-a)/n for some positive integer n. In simpler terms, the figures are partitioned into pieces, areas are calculated for each piece, and the summed pieces are the estimate of bias (Crocker & Algina, 1986; Henson, 1999).

For Lord's (1980) null hypothesis the formula is basically comparing the scores from each subpopulation by using the ICCs of item difficulty and probability of answering the item correctly. Lord (1980) also developed a similar test that is available for the two-parameter model, which will not be discussed in the present paper.

Insert Figure 1 about here

Figure 1 illustrates a three-parameter model. This ICC shows item bias against a specific group. For equal ability, the group whose ICC is above the other has a greater chance of getting the item correct. ICC for males lies consistently below the ICC for females, so the item is biased against males.

Insert Figure 2 about here

Figure 2 illustrates a three-parameter model with nonuniform bias. At low ability range, the item is biased against females, because the ICC for females lies below the ICC for males. At the high ability range, the item is biased against males, because the ICC for males lies below the ICC for females (Crocker & Algina, 1986; Fisk, 1991; Henson, 1999).

One-Parameter Rasch Model

For the one-parameter model, the b parameter for each item and person abilities (θ) are estimated. The item difficulty parameter is the sole estimate of performance. Discrimination is set at a constant (all items are presumed to be equally discriminating), and a chi square is used to determine fit (Cantrell, 1999; Henson, 1999).

Insert Figure 3 about here

Figure 3 illustrates a one-parameter Rasch Model ICC of an item where the ICCs are not the same for the two groups. The item is biased against blacks. The ICC for whites is consistently higher than the ICC for blacks. Also note that the bias is greatest for the middle ability range, because that is where the greatest difference is (Fisk, 1991; Lawson, 1991).

There are several important points to remember when using ICCs for item bias. Subpopulations can be unidimensional or multidimensional. If the ICCs vary and are unidimensional (one latent trait) in each group the ICCs variance may be due to measurement of different latent traits. Therefore, sources of variation are not equal. If the groups are multidimensional then item bias could result, yet one will NOT KNOW if the difference is due to item bias or multidimensionality (more than one latent trait in a group). Unfortunately, there is not an adequate procedure for measuring dimensionality;

therefore, results will always have some ambiguity. This ambiguity is left up to the researcher and his/her ability to think about the process and data at hand (Crocker & Algina, 1986; Fisk, 1991; Henson, 1999).

IRT is generally the best way to go when a large sample size and correct technology is available. If a large sample size or needed technology is unavailable there are other options, such as the chi square method (Henson, 1999).

Chi-square Method

Crocker and Aglina (1986) stated that Chi-square,

> essentially defines an item as unbiased if, within a group of examinees with scores in the same test score interval, the proportion of examinees responding correctly to the item is the same for both subpopulations. (p. 383)

For chi-square the observed score replaces the latent trait in IRT. The observed score scale is divided into several intervals, and within each interval the subpopulations are compared by proportions of correct item responses. Detection of item bias occurs if proportions vary across groups (Crocker & Algina, 1986). The following steps will illustrate the calculation of the proportion of those who answered the item correctly (P) for an interval within subgroups, and a proportion of those who answered the item correctly for an interval between subgroups. In reality, this will be done for each interval. For example, if there are 20 intervals, then the following process would occur 20 times.

First, the proportion (P) for subgroup intervals is calculated. There will be j amount of proportions for j amount of intervals for subgroup 1 (N1) and subgroup 2 (N2). O is the amount of items correct in the subgroup.

$$P_{1j} = O_{1j}/N_{1j}$$

The number of examinees that got the item correct in the examinees' subgroup is divided by the amount of examinees in the same subgroup to give the proportion of examinees in the first group and jth interval with the correct item.

Second, the proportion of the jth interval for both subgroups is determined.

$$P_{\cdot j} = \frac{O_{1j} + O_{2j}}{N_{1j} + N_{2j}}$$

The number of examinees that got the items correct in both subgroups are added and divided by the sum of the amount of examinees from both subgroups. Once all the P's are calculated, as shown in Table 2, chi-square can be calculated.

Insert Table 2 about here

The two statistical formulas used for chi-square are Camilli's statistic and Scheuneman's statistic. Scheuneman (1979) has a comprehensive review of this chi-square statistic, but Camilli's chi-square statistic will be the main focus in this paper. Scheuneman uses the follwowing formula:

$$\chi_C^2 = \sum_{j=1}^{J} \frac{N_{1j}N_{2j}(P_{1j} - P_{2j})^2}{(N_{1j} + N_{2j})P_{\cdot j}(1 - P_{\cdot j})} = \Sigma\chi_j^2$$

Camilli's chi-square is the sum of the chi-square for each interval of the subgroups. The formula incorporates both subgroups into the calculation for chi-square. Degrees of freedom for Camilli's chi-square is the number of intervals (J). The magnitude of this chi-square can be considered an index of the amount of bias in the item.

Limitations of Chi-square. There are several caveats to be aware of when using chi-square, and without taking note of the following the magnitude of the bias can be

influenced. First, a sufficient number of examines in each interval is important. There need to be incorrect scores within each interval, and at least ten to twenty intervals are needed. Also, frequencies of at least five correct responses are also beneficial.

Lastly, item bias may be an artifact of measurement error. When there is only control for the observed score differences, which is what is happening when scores are put into intervals, there is room for measurement error through true score differences. A subpopulation difference in item difficulty does not necessarily mean item bias; it could mean a difference in the construct of interest of the subpopulations.

Delta Plots

The delta plots method for detecting item bias is based on the item difficulty values (p), which was developed by Angoff (1972) and Angoff and Ford (1973). Items are deemed unbiased when item difficulties from group 1 are perfectly correlated with group 2, thus creating a straight line in the scatterplot. This is illustrated by placing item difficulties for group 2 on the y axis and item difficulties for group 1 on the x axis (Crocker & Algina, 1986). Item bias may be detected in a delta plot method when all items do not lie on a straight line.

Crocker and Algina (1986) defined a set of items as unbiased if the item difficulties for group 1 and group 2 are perfectly correlated. However, while a perfect correlation may exist between group 1 and group 2, item difficulties of a sample from each group may produce some degree of scatter.

A high correlation of item difficulties between groups is common if the rank order of difficulty is essentially the same for group 1 and group 2. Therefore, the goal of the delta plot method is to find possible outliers which contribute the most to item-by-group

interaction. Outliers are what the researcher is most interested in when studying item bias through the delta plot method (Fisk, 1991). Once outliers are determined, they are labeled as possible biased items.

When using the delta plot method there are three values that are necessary to illustrate possible item bias detection; item difficulty, z-scores, and delta measures. First, item difficulty scores (p-values) for the two different groups are computed on the items chosen. Second, z-scores are found by using a z-score table and finding the cut off score for the p-value of each item. Third, the cut off z-scores for the p-values are then converted to a normal deviate with an arbitrary mean and standard deviation. In the following illustration, the values 13 (mean) and 4 (standard deviation) are used. Deltas, the transformed normal deviates, are calculated using the following formula, and then plotted on a bivariate graph to illustrate possible biased items (Fisk, 1991):

$$\Delta = 4z + 13$$

Table 4 shows the three values discussed earlier that are needed to create the delta plot. The values were calculated from a hypothetical study of two groups of examinees on a 15-item, dichotomously-scored test.

Insert Table 3 about here

The values used to create the delta plot are the delta values in column 6 (d1) and column 7 (d2) in Table 4. The delta plot is illustrated in Figure 3.

Insert Figure 3 about here

The items deviated from the line are seen as possibly biased items. The point that is shaped like a box in Figure 3 would be considered for possible item bias because the

box deviates from the line. Unfortunately, there is not a set cut-off value to detect when the item is actually biased. This is left up to the researcher to decide.

A more formal method given by Angoff (1982), and discussed by Crocker and Algina (1986) and Fisk (1991), measures the distance of each outlier from the major axis of the ellipse formed by the scatter plot of the data. This is discussed in detail in Fisk (1991).

The main area of concern for the delta-plot method is that the distribution of the ability of the examinees influence the results of the analysis (Fisk, 1991; Ironson, 1982). Angoff (1982) also provided a list that charts the advantages and disadvantages of the delta plot method (Fisk, 1991).

Another problem of the delta plot method is that even though an entire set of items may be unbiased according to the latent trait definition, and the delta plot method may nevertheless indicate bias. This is due to the item discrimination parameters being unequal, and items with large item discriminations tend to appear more biased (Crocker & Algina, 1986). Crocker and Algina (1986) presumed an extensive discussion of how to resolve many of the drawbacks to the delta plot method.

Conclusions

The purposes of item bias analyses are to investigate whether test scores are affected by different sources of variance in the various subpopulations, and if different sources of variance are found, to determine if an unfair advantage exists. The identification of item bias is a challenge, and the methodology chosen to test for item bias is yet another challenge. When choosing the best item bias detection methodology, the researcher must consider the application of the results, the availability of software

needed, and the practicality of implementing the methodology chosen (Fisk, 1991). The IRT model leads the other item bias methods with being the most theoretically sound, yet statistically complex procedure. Chi-square and delta plot methods are not as theoretically sophisticated, yet much more practical and easier to implement.

References

Angoff, W. H. (1982). Uses of difficulty and discrimination indices for detecting item bias. In R. A. Berk (Ed.), *Handbook of methods for detecting item bias.* Baltimore: Johns Hopkins University Press.

Angoff, W. H. & Ford, S. F. (1973). Item race interaction on a test of scholastic aptitude. *Journal of Educational Measurement, 10*, 95-105.

Cantrell, C. E. (1999). Item Response Theory: Understanding the one-parameter Rasch model. In B. Thompson (Ed.), *Advances in social science methodology* (Vol. 5, pp. 171-192). Stamford, CT: JAI Press.

Crocker, L., & Algina, J. (1986). *Introduction to classical and modern test theory.* New York: Holt, Rinehart and Winston.

Fan, X. (1998). Item Response Theory and classical test theory: An empirical comparison of their item/person statistics. *Educational and Psychological Measurement, 58*, 357-381.

Fisk, Y. (1991, January). *A brief overview of three methods for detecting item bias.* Paper presented at the annual meeting of the Southwest Educational Research Association, San Antonio. (ERIC Document Reproduction Service No. ED 328 578)

Henard, D. H. (2000). Item Response Theory. In L. Grimm & P. Yarnold (Eds.), *Reading and understanding more multivariate statistics* (pp. 67-98). Washington, DC: American Psychological Association.

Henson, R. K. (1999, January). *Understanding the one-parameter Rasch model of Item Response Theory.* Paper presented at the annual meeting of the Southwest

DC: American Psychological Association.

Henson, R. K. (1999, January). *Understanding the one-parameter Rasch model of Item Response Theory.* Paper presented at the annual meeting of the Southwest Educational Research Association, San Antonio. (ERIC Document Reproduction Service No. ED 428 078)

Jenson, A. R. (1980). *Bias in mental testing.* New York: The Free Press.

Lawson, S. (1991). One parameter latent trait measurement: Do the results justify the effort? In B. Thompson (Ed.), *Advances in educational research: Substantive findings, methodological developments* (Vol. 1, pp. 159-168). Greenwich, CT: JAI Press.

McKinley, R., & Mills, C. (1989). Item response theory: Advances in achievement and attitude measurement. In B. Thompson (Ed.), *Advances in social science methodology* (Vol. 1, pp. 71-135). Greenwich, CT: JAI Press.

Scheuneman, J. (1979). A new method for assessing bias in test items. *Journal of Educational Measurement, 16*, 143-152.

Table 1

Area Measure (A_g) of Item Bias

Sex

	Male		Female		
Item	b1g	a1g	b2g	a2g	Ag
1	-2.753	.672	-1.467	1.98	1.13
2	.130	1.252	.222	.723	.47
3	- .959	3.167	-1.602	.989	.77
4	-2.46	.531	-1.734	.592	.32
5	-1.865	.715	-3.202	.5891	1.02
6	-.875	2.167	-1.278	.909	.60
7	-.283	1.103	.139	.718	.52
8	-.733	1.162	-.428	.963	.31
9	-.531	2.405	-.624	.903	.55
10	.155	1.546	.115	.686	.64

Table 2

TABLE 16.3. Illustrative Data for Calculating χ_c^2 and χ_s^2

Interval	Score Level	N_{1j}	O_{1j}	P_{1j}	N_{2j}	O_{2j}	P_{2j}	$P_{\cdot j}$
1	13–14	25	22	.880	315	300	.952	.947
2	12	24	18	.750	110	99	.900	.873
3	10–11	48	23	.479	118	93	.788	.698
4	1–9	65	14	.215	92	33	.358	.299

From J. Scheuneman, A new method for assessing bias in test items, *Journal of Educational Measurement, 16,* 143–152. Copyright 1979 by the National Council on Measurement in Education, Washington, D.C. Adapted by permission.

Table 3

Item Difficulty (p), z-scores (z), and Delta Values (d) For Group 1 and Group 2

Item	p1	p2	z of p1	z of p2	d1	d2
1	.529412	.529412	0.073792	0.073792	13.29517	13.29517
2	.705882	.529412	0.541395	0.073792	15.16558	13.29517
3	.352941	.294118	-0.37739	-0.54139	11.49043	10.83442
4	.647059	.647059	0.377393	0.377393	14.50957	14.50957
5	.529412	.529412	0.073792	0.073792	13.29517	13.29517
6	.058824	.117647	-1.56473	-1.18683	6.741094	8.252674
7	.235294	.235294	-0.72152	-0.72152	10.11391	10.11391
8	.352941	.294118	-0.37739	-0.54139	11.49043	10.83442
9	.705882	.647059	0.541395	0.377393	15.16558	14.50957
10	.058824	.058824	-1.56473	-1.56473	6.741094	6.741094
11	.352941	.705882	-0.37739	0.541395	11.49043	15.16558
12	.647059	.588235	0.377393	0.223008	14.50957	13.89203
13	.470588	.470588	-0.07379	-0.07379	12.70483	12.70483
14	.411765	.352941	-0.22301	-0.37739	12.10797	11.49043
15	.058824	.058824	-1.56473	-1.56473	6.741094	6.741094

Figure 1

Figure 1

Hypothetical ICCs for a 3-paramenter model depicting item bias

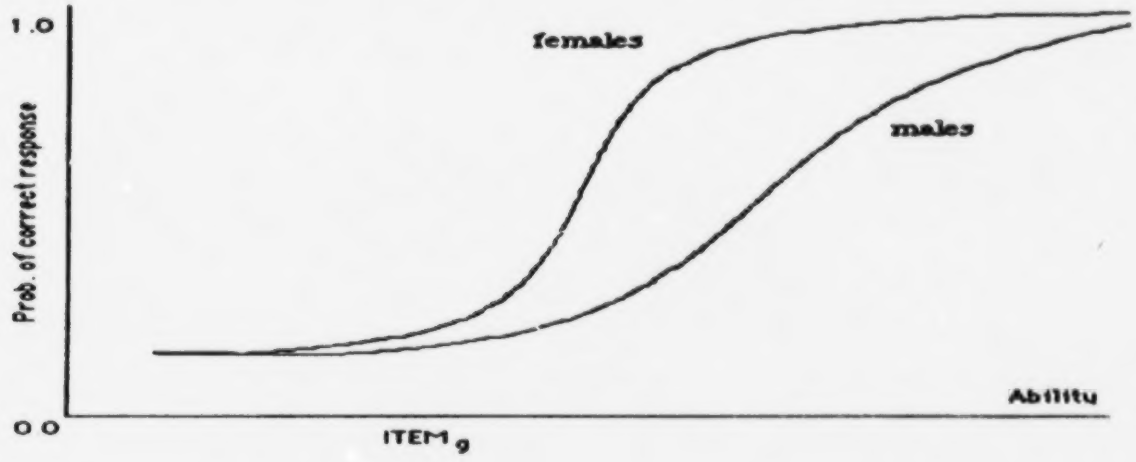

($P_g(\theta)$ = probability of a correct response on item g as a function of ability)

Figure 2

Figure 3

Hypothetical ICCs for a 1-parameter model depicting biased item

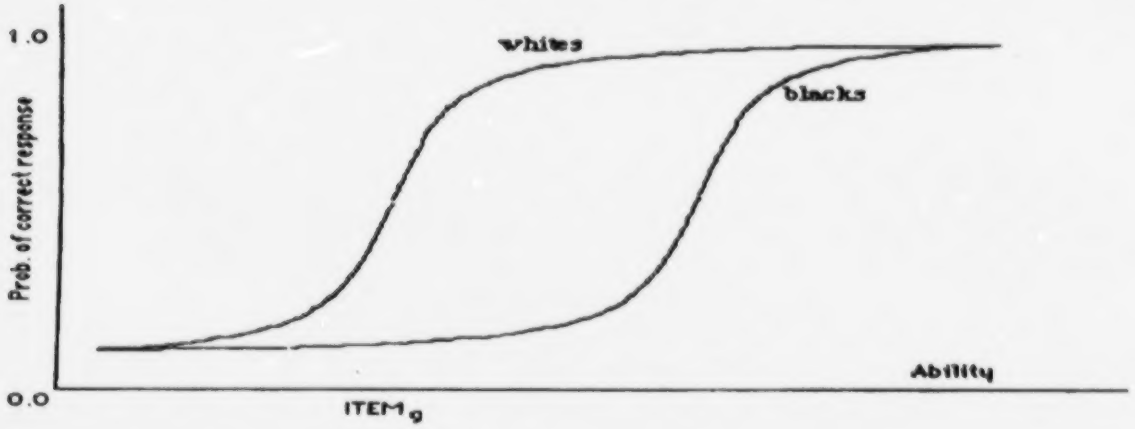

($P_g(\theta)$ = probability of a correct response on item g as a function of ability)

Figure 3

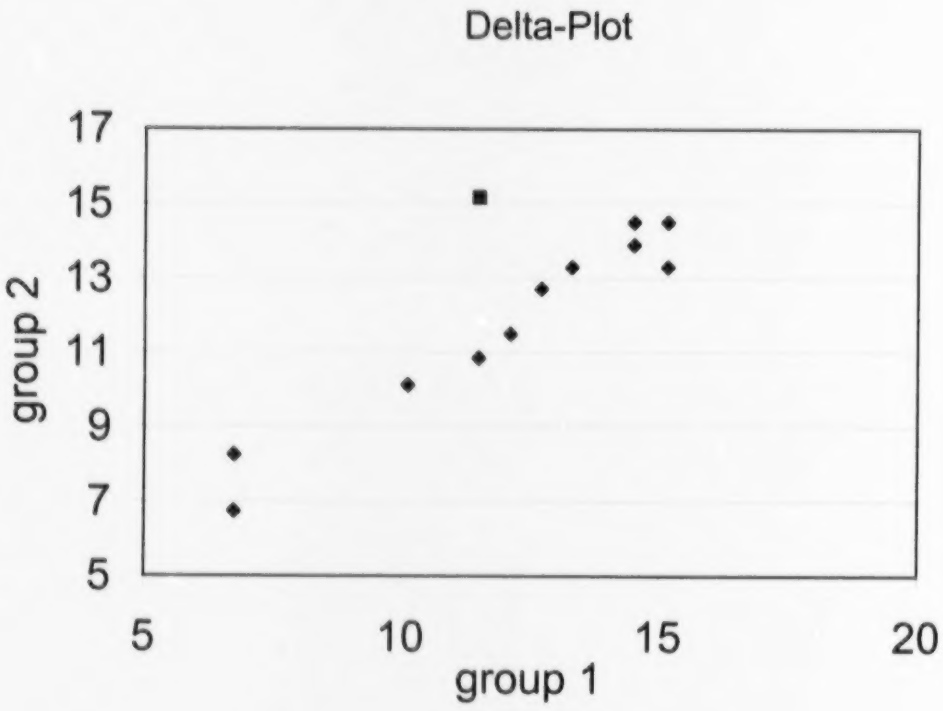

III. DOCUMENT AVAILABILITY INFORMATION (FROM NON-ERIC SOURCE):

If permission to reproduce is not granted to ERIC, or, if you wish ERIC to cite the availability of the document from another source, please provide the following information regarding the availability of the document. (ERIC will not announce a document unless it is publicly available, and a dependable source can be specified. Contributors should also be aware that ERIC selection criteria are significantly more stringent for documents that cannot be made available through EDRS.)

Publisher/Distributor:
Address:
Price:

IV. REFERRAL OF ERIC TO COPYRIGHT/REPRODUCTION RIGHTS HOLDER:

If the right to grant this reproduction release is held by someone other than the addressee, please provide the appropriate name and address:

Name:
Address:

V. WHERE TO SEND THIS FORM:

Send this form to the following ERIC Clearinghouse:

University of Maryland
ERIC Clearinghouse on Assessment and Evaluation
1129 Shriver Laboratory
College Park, MD 20742
Attn: Acquisitions

However, if solicited by the ERIC Facility, or if making an unsolicited contribution to ERIC, return this form (and the document being contributed) to:

ERIC Processing and Reference Facility
1100 West Street, 2nd Floor
Laurel, Maryland 20707-3598

Telephone: 301-497-4080
Toll Free: 800-799-3742
FAX: 301-953-0263
e-mail: ericfac@inet.ed.gov
WWW: http://ericfac.piccard.csc.com

EFF-088 (Rev. 9/97)
PREVIOUS VERSIONS OF THIS FORM ARE OBSOLETE.